Microsoft

Windows NT® Server

Server Operating System

Windows NT 5.0 Namespace Design

White Paper

Abstract

This paper describes the Microsoft® Windows NT® 5.0 operating system Active Directory namespace architecture, including the forest and tree domain structure, organizational units, the global catalog, trust relationships, and replication. It then provides examples of namespace implementations and describes the architectural criteria that network architects and administrators should consider when designing an Active Directory namespace for the Enterprise.

© 1997 Microsoft Corporation. All rights reserved.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This White Paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT.

Microsoft, the BackOffice logo, Windows, and Windows NT are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Other product or company names mentioned herein may be the trademarks of their respective owners.

Microsoft Corporation • One Microsoft Way • Redmond, WA 98052-6399 • USA

0997

CONTENTS

INTRODUCTION.....1

ACTIVE DIRECTORY FEATURES.....2
Domains 2
Forests and Trees 2
Summary of Tree and Forest Features 3
Tree and Forest Details 4
Organizational Units (OUs) / Containers 6
Global Catalog 6
Trust Relationships 7
Replication Boundaries 7

NAMESPACE SCENARIO CONSIDERATIONS.....8
Separate Internal and External Namespaces 8
Same Internal and External Namespace 9

NAMESPACE SCENARIO COMPARISON.....11
Separate Internal and External Namespaces 11
Advantages 11
Disadvantages 11
Same Internal and External Namespace 11
Advantages 11
Disadvantages 11

DEFINING THE NAMESPACE ARCHITECTURE.....13
Domain Tree 13
Root Domain 13
First-Layer Domains 13
Second-Layer (Child) Domains 14
Partner/Joint Venture Domains 15
Organizational OUs 15

FOR MORE INFORMATION.....17

APPENDIX A – CORPORATE WINDOWS NT 5.0 ACTIVE DIRECTORY NAMESPACE.....18

APPENDIX B – SAMPLE TOP-LEVEL OU STRUCTURES.....19

APPENDIX C – ISO 3166 TWO-CHARACTER COUNTRY CODES.....20

APPENDIX D – TWO-LETTER US POSTAL CODES.....22

ACTIVE DIRECTORY FEATURES

The Microsoft® Windows NT® operating system version 5.0 includes a newly designed directory service called the *Active Directory*. The Active Directory extends the features of previous Windows® operating system-based directory services and adds entirely new features. The Active Directory is *secure, distributed*, *partitioned*, and *replicated*. It is designed to work well in any size installation, from a single server with a few hundred objects to thousands of servers and millions of objects. The Active Directory adds many new features that make it easy to navigate and manage large amounts of information, generating time savings for both administrators and end users. These features include a new domain model and hierarchical namespace.

This paper describes the Windows NT 5.0 Active Directory namespace, including the forest and tree domain structure, organizational units, the global catalog, trust relationships, and replication. It then provides examples of namespace implementations and describes the architectural criteria that network architects and administrators should consider when designing an Active Directory namespace for the Enterprise. By following the recommendations in this paper, the Enterprise network architect should be able to design a namespace that is capable of withstanding company reorganizations without expensive restructuring.

The following sections describe the features of the Windows NT 5.0 Active Directory and compare them to the corresponding features in Windows NT *3.x*/4.0.

Domains

Just as it was in the Microsoft Windows NT *3.x*/4.0 architecture, the core unit in the Microsoft Windows NT 5.0 Active Directory is the domain. All network objects[1] exist within a domain. Access to objects is controlled by Access Control Lists (ACLs) populated with Access Control Entries (ACEs). Within a domain, access permissions are cumulative unless explicitly denied. Administration rights are limited to domain boundaries by default.

Domains are also units of replication. A single domain can span multiple physical locations or sites. Unlike the single-master model used by Windows NT *3.x* and 4.0, with its Primary Domain Controllers (PDCs) and Backup Domain Controllers (BDCs), the Active Directory uses a multimaster “peer controller” model. Thus, all of the domain controllers (DCs) that are authoritative for a given domain can receive changes

[1] Directory service objects are not true “objects.” They do not contain extensible methods. Instead, they are attributed representations of various resources that exist on the network.

directly and propagate those changes. This allows intersite replication to occur within a domain, even if any single DC is down.

Multiple domains within the Active Directory are linked into a *domain tree*. Users in the linked domains can access resources in other domains via transitive trust relationships that exist hierarchically among all of the domains in the domain tree. However, administrative rights are not inherently transitive. Thus, an additional layer of security is added by limiting the scope of domain administrative rights.

Forests and Trees

As stated above, Windows NT 5.0 introduces the concept of a *tree*, a hierarchical structure of domains that form a contiguous namespace and share a common schema, configuration, and global catalog, with Kerberos trust among all members of the tree. A *forest* is a collection of one or more trees connected by Kerberos trusts organized as peers. This provides companies with the option of constructing their enterprise from separate, distinct disjoint namespaces-

Trees and *forests* are a refinement of the original domain tree concept and are designed to provide a multidomain structure in which:

- Naming and security relationships among member domains are clear and intuitive.
- Security principals have a preferred logon name that is easily determined by their position in the structure.
- Behavior of domain move and rename functions are clear and intuitive.
- Terminology is unambiguous and easily grasped.
- Both DC= and “classic X.500” names are supported.

The net effect of revising the domain tree concept as described in this document is a Windows NT 5.0 product that is easier to understand, easier to use, and more competitive. The difference in presentation between the original domain tree concept and a Forest of Trees is shown in Figures 1 and 2 below.

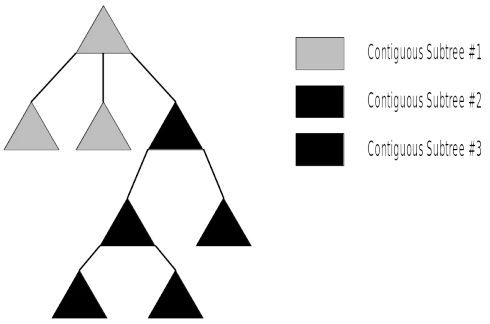

Figure 1. Domain Tree per Existing Implementation

Figure 1 depicts an enterprise consisting of three subtrees, each internally contiguous, joined into a single *discontiguous* domain tree. Figure 2 shows the same enterprise implemented as a *forest* in which each tree of the enterprise is distinct.

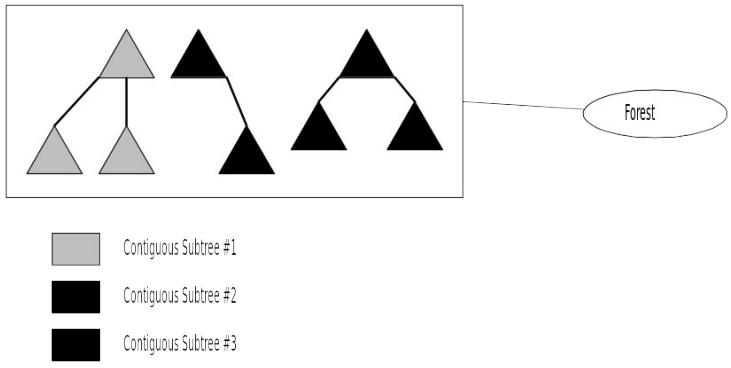

Figure 2. Forest Per Proposed Implementation

Summary of Tree and Forest Features

- ***Tree***—a group of Windows NT domains that form a contiguous namespace, sharing a common schema, configuration, and global catalog, with Kerberos trust among all members of the tree.
- ***Forest***—a set of trees that share a common schema, configuration, and global catalog, with Kerberos trust among all members of the forest. A forest does *not* form a contiguous namespace.
- The trust hierarchy within a tree follows the naming hierarchy; this is both intuitive and easy to implement.
- The trust hierarchy in a forest is an arbitrary graph. By default the trust hierarchy is formed in the order the trees are joined to the forest; that is, the “top” is the root domain of the first tree and subsequent trees trust the “top.” The user does not know or care what the trust tree looks like. The trust hierarchy is exposed to the

administrator for management purposes, but is not otherwise significant.

- Trust within a tree or forest can be administratively adjusted to optimize Kerberos ticket-granting referrals.
- The top of the browsing user interface is a list of trees.
- Enterprises can be single-tree or multitree. Naming within a given tree is *always* contiguous.
- A deep search from the root of a domain flows down throughout a tree by automatically searching through subordinate references.
- Search behavior in a forest is intuitive; a deep search is restricted to the tree where the search was initiated. A search from the root of the namespace searches the global catalog. This makes determining which referrals to search relatively easy. Referrals within a tree are always searched, but referrals to other trees in a forest are never searched.
- Renaming a domain in a tree is an implicit rename of all children of the renamed domain; for example, renaming a.b.com to a.newname.com implicitly renames child.a.b.com to child.a.newname.com.
- A user logs on with a user principal name (UPN). A UPN is the user's Security Account Manager (SAM) account name composed with the user’s tree name. A user is always a member of exactly one tree. If the user changes trees, the UPN changes to the new tree.
- The name of a tree is the Domain Name System (DNS) name of the root domain. Renaming the root domain renames the tree and (as stated above) all child domains in the tree.

Tree and Forest Details

A *tree* is a set of one or more Windows NT 5.0 domains sharing a common schema, configuration, and global catalog, joined together to form a **contiguous** namespace. All domains in a given tree trust each other through transitive hierarchical Kerberos trust relationships. The minimal tree is a single Windows NT domain. A larger tree can be constructed by joining additional domains as children to form a larger contiguous namespace.

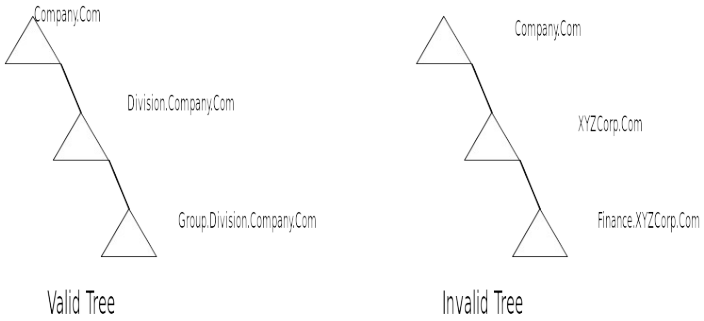

Figure 3. Trees Are Always Contiguous

A tree has a name. The tree name is *always* the DNS name of the domain at the root of the tree. Children of the tree root are *always* contiguous with it in the namespace. The DNS names of the children of a tree root reflect this; that is, the children of *Somedomain.com* are always children of *Somedomain.com* in the DNS namespace—for example, *child1.somedomain.com*, *child2.somedomain.com*, and so forth. A tree may use either the DC= convention for naming or the classic X.500 style (OU=child1,O=SomeOrg,C=US). However, a single tree cannot support both styles.

A *forest* is a set of one or more trees that *do not* form a contiguous namespace. All trees in a forest share a common schema, configuration, and global catalog. All trees in a forest trust each other through transitive, hierarchical Kerberos trust relationships. Unlike trees, a forest does not need a distinct name. A forest exists as a set of cross-reference objects and Kerberos trust relationships known to the member trees. Trees in a forest form a hierarchy for the purposes of Kerberos trust; the tree name of at the root of the trust tree can be used to refer to a given forest.

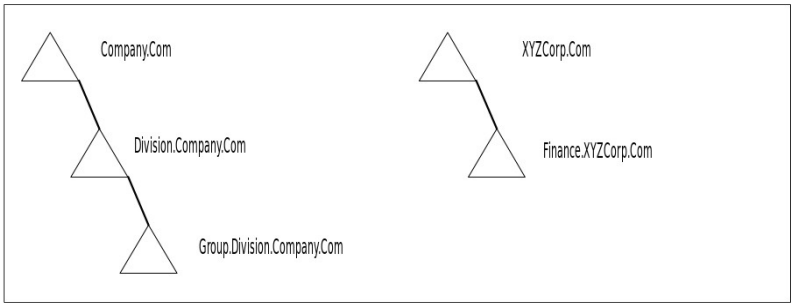

Figure 4. Multiple Trees in a Forest

The major difference between a forest and a domain tree is the removal

of the requirement that the member domains always form a tree. Disjointed subtrees were joined together arbitrarily to form a domain tree. The behavior of this arbitrary tree is not intuitive and has made the entire concept awkward to use and describe.

For example, in the arbitrary tree, the search behavior is not intuitive. Search operations are only propagated if the child is contiguous with the parent, and ACLs are not propagated at all. There is no obvious way to organize such a tree, and the advantages of one particular structure over another are not clear.

In the forest approach, the component trees are peers. The relationship between the peer trees is simply one of Kerberos trust. Trees are trees of domains that form a contiguous namespace. The behavior of the tree is intuitive; search operations flow down the tree.

In a forest of trees, the structure reflects the functionality. The contiguous namespaces are trees; these are distinct and behave intuitively. If there are multiple disjointed trees in the enterprise, they are peer members of the forest. The trees in a forest are not represented as a larger tree, eliminating the confusing situation caused by forcing disjointed namespaces into an arbitrary tree structure.

Organizational Units (OUs)/Containers

Within a domain, directory service objects are organized into *organizational units* (OUs). It is important to note that Windows NT 5.0 OUs are not equivalent to X.500 OUs. Instead, Windows NT OUs are administrative boundaries. They are used for organizing user and resource objects into logical administrative groups. Various administrative tasks (such as access rights specification) can then be delegated to the administrator for a specific OU, thereby freeing domain administrators from having to support such changes by proxy.

OUs also provide inheritance of access rights. This allows access to resources specific to a particular organization to be restricted to members of that OU.

Global Catalog

Today with the Microsoft Windows NT *3.x*/4.0 domain model, as the number of resource domains increases, users are presented with numerous domains in the browse listing. To simplify the user’s view, Windows NT 5.0 provides a search mechanism called the *global catalog*.

The global catalog (GC) is a partial index of select objects in the domain tree, combined with a search engine. To find a resource in the domain tree, a user queries the GC for that resource, based on one or more known attributes of the target resource. The GC returns the location of

the desired resource. One of the object attributes included in the GC is the Access Rights that pertain to that object. This reduces the obvious security risk posed by the GC. If a user does not have rights to access an object, the user's query will fail. Thus, the GC does not provide a single point of risk exposure from a security perspective

Trust Relationships

As previously mentioned, a transitive trust relationship exists among the domains in an Active Directory domain tree. This eliminates the need to explicitly maintain and manage two-way trusts between all of the domains on a corporate network. The root domain of the domain tree is called the root. All other domains in the domain tree are hierarchically branches from the root.

The two trust mechanisms are transitive Kerberos trusts, and explicit Windows NT 3.1/4.0-style one-way trust relationships. The first mechanism allows users to gain access to resources anywhere throughout the organization (subject to access control), without setting up explicit trusts to each of the other participating domains. The second mechanism allows for domains that are either members of another Active Directory domain tree or that do not support Windows NT Active Directory authentication to have limited access to the domain tree. The one-way trust relationship limits the scope of the authenticated access to the member domain that is explicitly trusted. This provides a mechanism for limiting the network resources that are visible and available to members of this partially trusted domain.

Replication Boundaries

The Windows NT Active Directory uses multimaster replication. Domains are units of replication. Any change at one site or another within a domain is replicated to the other sites. (A *site* is defined as one or more well-connected network segments.) Since there is no one master domain —or even a Primary Domain Controller—changes can be made simultaneously at all of the various sites and/or controllers within a domain. The Active Directory uses update sequence numbers (USNs) to track changes on a per-attribute basis. The USNs and timestamps of the changes are used to resolve conflicting changes.

Certain changes—such as the addition or deletion of a domain and changes in the domain object schema—have consequences for the entire domain tree. These operations require some lockout mechanism to ensure that the change is propagated prior to the time that the next domain-wide change is begun. The Active Directory addresses this lockout requirement by designating a DC within the enterprise as the *floating single master* (FSM). Within the organization, only one DC can

serve as the FSM at any point in time. At different times, different DCs within the domain can assume this role. In the event that the FSM fails, another DC can be promoted to take on this role.

This section describes different approaches to implementing the Active Directory namespace structure. The next section, “Namespace Scenario Comparison,” discusses the advantages and disadvantages of the different approaches.

Separate Internal and External Namespaces

In this example, the company uses separate internal and external namespaces. *Company.com* is the name used outside the firewall, and *cpny.com* is the name used inside the firewall. Figure 5 illustrates this implementation:

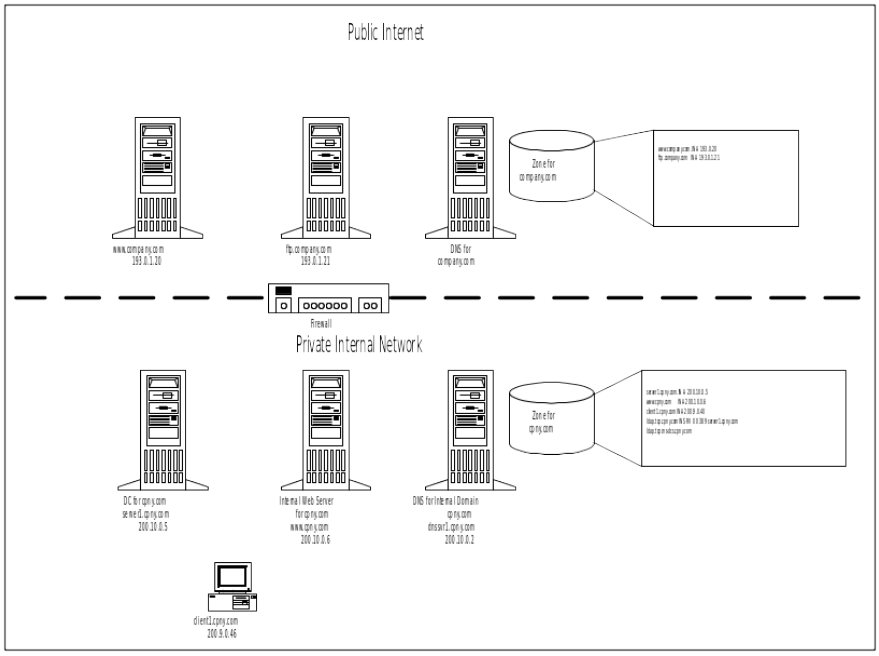

Figure 5. Different Internal and External Namespaces

In this configuration, there are separate names used inside and outside the corporation. As illustrated in Figure 5, company.com is the external namespace presented to the Internet community. Inside the company, cpny.com is used. This architecture requires that two namespaces—company.com and cpny.com—be reserved with an Internet DNS registration authority. The internal name, cpny.com, should be reserved to prevent that name from being used on the public Internet. Failure to reserve this name would prevent internal clients from accessing this namespace on the public

Internet in the future, because the client would be unable to distinguish between the internally implemented name and the publicly assigned name.

The DNS zone configuration follows a traditional format, with a zone setup to resolve resources on the public Internet for company.com, and a second zone defined inside the firewall that contains the resources available to clients on the company internal network for cpny.com. A client is able to clearly distinguish between an internal resource and an external one based upon the fully qualified domain name.

Same Internal and External Namespace

In this example, the company uses the same name for the internal and external namespaces. *Company.com* is used both inside and outside the company. This implementation has the following requirements:

- Clients on the internal network must be able to access all of the servers on the inside of the firewall in addition to the servers that reside outside the firewall on the Internet.
- Clients accessing the site from outside the firewall must not be able to access or resolve names of internal resources to protect company resources from attack.

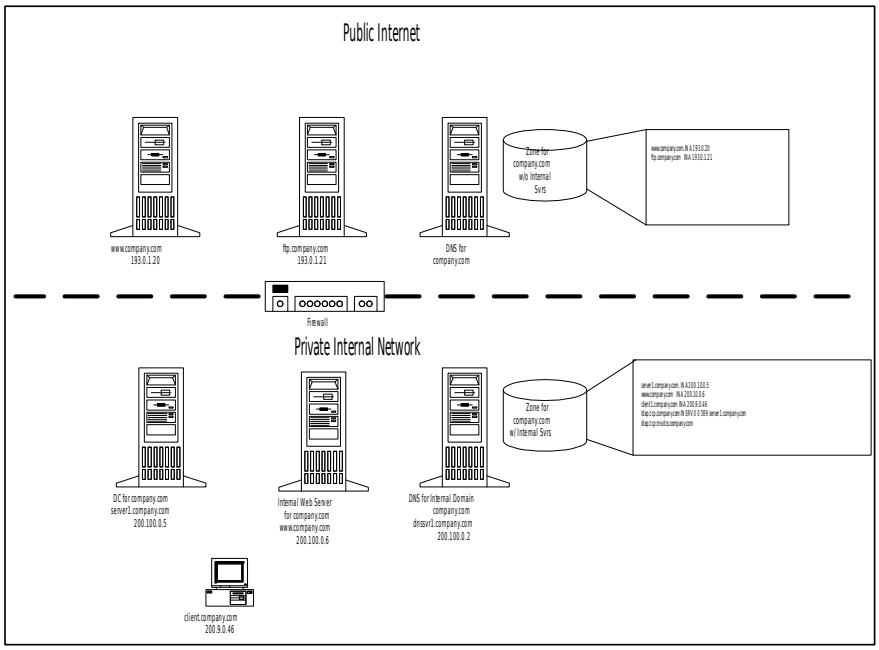

Figure 6. Same Namespace both Internally and Externally

In this configuration, two separate zones for company.com exist as shown in Figure 6. One zone will be created that provides name resolution for resources outside the firewall. The zone file residing on the DNS servers outside the firewall references publicly accessible resources

such as Web servers, FTP servers, and mail servers. This configuration prevents clients external to the corporation from querying DNS for internal-only corporate servers. It should also be noted that no directory services servers reside outside the firewall, so no Kerberos Distribution Centers are more susceptible to attack.

The zone file created inside the company network contains a superset of servers available to internal-only company resources. This enables clients within the corporation to resolve all company resources.

One challenge presented by using the same namespace both internally and externally is enabling internal clients to access publicly available resources. One way to ensure corporate network clients have access to Internet servers is to mirror those services inside the corporate network. Proxy client software should be configured to treat names ending in *.company.com as internal to the corporate network, and mirrors of the applicable services would be entered in the internal company.com DNS zone. Mirroring has the useful side effect of making access to company Internet services from the corporate network more efficient. The following are additional proxy client configuration considerations:

- Exclusion lists—clients can be configured such that requests for access to servers residing on the public Internet are not passed outside the firewall, but rather redirected to an internal mirror for the resource.
- Proxy client browsers receive a different view of network resources from inside the company versus outside the company. Users must be made aware of the fact that the view of corporate network resources will differ depending on the vantage point of the client. Users must understand that when they access company.com from the public Internet, they will not be able to view internal corporate network resources for security purposes.

Additional client browser configuration considerations will be addressed in a subsequent White Paper.

The following paragraphs describe the advantages and disadvantages of the Active Directory namespace implementations that were described in the previous section.

Separate Internal and External Namespaces

The situation in which the namespace *company.com* is used externally and *cpny.com* is used internally requires that separate names exist on either side of the firewall and resources exist inside the company network. This approach has the following advantages and disadvantages.

Advantages

- A clear distinction exists between resources that are internal and those that are external.
- The structure is easier to manage from an operations perspective because the disjointed namespaces can be managed separately and have no overlap.
- The client browser configuration is simpler (an exception list blocks all names that end in *.cpny.com from going outside the firewall).
- The proxy client configuration is simpler (an exception list identifies all names ending in *.cpny.com as internal).

Disadvantages

- A user's logon name is different from his or her e-mail name (for example, Jane Smith would use the logon jsmith@cpny.com and the e-mail address jsmith@company.com. This increases user education requirements and adds a potential level of confusion.
- This approach requires that the company reserve multiple names with an Internet DNS registration authority. This potentially poses a problem if someone else currently owns the selected name and is unwilling to concede the name.

Same Internal and External Namespace

The situation in which the namespace *company.com* is used both externally and internally requires that the same name exist on both sides of the firewall. This approach has the following advantages and disadvantages.

Advantages

- The tree name, company.com, is consistent on the public Internet and private intranet.
- The same logon and e-mail names can be used (for example, the name jsmith@company.com would function as both the logon and e-mail ID).

Disadvantages

- This approach requires complex proxy client configurations—a proxy client must understand when to send a request outside the corporate network.
- This causes confusion in operations—administrators must understand the differences between publishing resources on the Internet and the intranet to prevent accidental publishing of internal-only content.
- Overlapping internal/external namespaces require synchronized management and present increased potential for confusion.
- Users must know that they will get a different view, depending on whether they are looking at the external or internal network.

DEFINING THE NAMESPACE ARCHITECTURE

When designing an Active Directory namespace architecture, you should consider the following design criteria:

- Replication traffic requirements.
- The need to accommodate organizational changes without expensive domain changes.
- The ability of the design to evolve as organizational needs change.

The following sections address these criteria, suggest one possible approach *(an organization is not forced into having a single domain tree as we describe below*), provide examples, and discuss the rationale behind the design approach taken.

Domain Tree

The recommended corporate network directory service architecture is a mixture of both Windows NT 5.0 domains and organizational units. In this architecture, Windows NT 5.0 domains are organized into a root domain, first-layer domains, and second-layer domains. Structuring the company in this manner provides more granular replication, the ability to limit the scope of domain administrators, and the ability to limit the number of global domain administrators.

Root Domain

The top domain of the namespace hierarchy represents the whole organization and is known as the *root* domain. It is from this one internal domain that all other domains originate. Only one name should be chosen for the root domain, and it should be easily recognizable. This reduces the costs associated with purchasing and maintaining another DNS name from an Internet DNS registration authority. It also provides consistency for user logon names and the company’s well-known corporate Internet e-mail names. The internal root domain is an empty Active Directory domain used solely for the directory service namespace. No directory service-enabled servers from the root domain will reside on the external public Internet.

First-Layer Domains

The first layer of domains under the root domain represents continental and geopolitical boundaries. The domains within this layer are created to minimize directory service replication and to provide a means for creating persistent names. The objective here is to create domains whose names don’t change. Changes in the overall domain architecture, such as domain collapses and recreation, is a difficult and potentially IT-intensive support proposition. A good namespace design should be capable of withstanding company reorganizations without the need to restructure

the existing domain hierarchy. First-layer domains contain the OUs of remote offices that have neither the hardware nor support resources to maintain a separate domain. By defining these remote offices as OUs, an organizational hierarchy is provided that accommodates promoting them to domains in the future, if such a need arises.

The naming convention for domains within this layer must be at least three characters long so that they do not conflict with the ISO 3166 two-character country code used to name second-layer domains and OUs. This first layer of domains should be relatively stable and unchanging. If additional domains are needed, they should be created as child domains under one of the existing first layer domains.

One possible naming convention for first-layer domains is shown in the following table.

Domain	Definition
CORPIT	Company IT Headquarters
NOAMER	United States of America and Canada
SOAMER	Mexico, Central America, and South America
NOPAC	Hong Kong and sites north of Hong Kong (Hong Kong, China, Korea, Taiwan)
SOPAC	Sites south of Hong Kong, including the India subcontinent over to but not including Afghanistan
EUROPE	Austria, Belgium, Czechoslovakia, Denmark, Finland, Greece, Croatia, Hungary, Ireland, Italy, Holland, Norway, Poland, Portugal, Romania, Russia, Spain, Slovakia, Slovenia, Sweden, Switzerland
MEAST	Israel, Saudi Arabia, Turkey, United Arab Emirates
AFRICA	Africa
PARTNERS	Business partners and companies to which work is outsourced
JVT	Joint ventures

Second-Layer (Child) Domains

The larger branch offices and sites are defined as second-layer or child domains, and are branched off of their corresponding first-layer (parent) domains. Child domains at the second layer should be countries only. Additional child (of child) domains can be created, if needed, off of the second-layer domains to support a specific location within a country.

The naming convention for domains and OUs are the same. This allows an OU to be promoted to a domain with as little user impact as possible. The two-character ISO 3166 standard is used for all locations except

locations within the United States. These locations do not have an ISO country code and therefore follow the two-letter U.S. postal code. The single exception to this convention is California, whose two-character postal code conflicts with Canada's ISO code; therefore, the elongated state abbreviation of "calif" is used. For ease of use, all names under the first layer are lower case. Use the following criteria to determine which remote offices would be child domains:

1. Remote offices with two or more data servers.
2. Remote offices with local technical IT support personnel.
3. Remote offices with staff of more than 1,000 people.
4. Remote offices that have sufficient network bandwidth and replication capabilities.

Partner/Joint Venture Domains

Partner and Joint Venture domains are owned and managed by your organization. They contain the user accounts and shared resources that are needed by specific partners and joint ventures. The specific user accounts of the partner's domain reside in disjointed domains that are linked to your corporate network directory service by explicit one-way trusts. The joint-venture domain also contains user accounts, but will be linked to the corporate network by Windows NT 5.0 transitive trusts. The explicit one-way trust links are not transitive across the domain tree; access is restricted to the network resources contained within the specified domain. This allows access to sufficient resources to do the necessary work without exposing all of the company's network resources to nonemployees.

Partners are defined as those companies that have explicit contracts with your company to provide services, and that require network connectivity to provide those resources. Individual vendors or contractors are more appropriately managed within the OUs for which they provide services. This allows for more careful tailoring of access rights for these individuals.

Joint Ventures are defined as those companies with which your company has joint venture or similar development agreements and which require connectivity to your corporate network. This access can also be provided for individual users by issuing X.509 certificates with appropriately defined expiration dates. The second approach adds more administrative attention.

Organizational OUs

Within various domains it is desirable to further organize resources and

users to reflect the business organization. OUs are the most appropriate vehicle for doing this. Since OUs can contain further OUs, the level of granularity can be extended as far as necessary. Changes within and between OUs can be administered relatively simply. Much of this administration can be delegated to the “owners” of the OUs. Just as it is not practical for a central IT organization to directly support all of the organizational needs of various departments, so too it is not practical for the central IT organization to administer all of the OUs that business-driven organization may require. It is therefore recommended that the central IT organization only directly support and define the first level of the business model organizational OUs. Since there is no replication or hardware costs associated with being an OU, the first layer of business organization OUs should be standard throughout the organization regardless of whether a particular location requires them or not. This provides consistency for company-wide support.

The following list provides examples of OUs and their respective departments. This list is provided as an example, and may need to be varied depending on specific organizational needs.

OU	Department
Legal	Legal
Fac	Facilities
Fin	Finance
Mktg	Marketing
Support	Support
Mfg	Manufacturing
Dist	Distribution
Hr	Human Resources
Sales	Sales
It	Information Technology
Exec	Executive Staff

Further organizational breakdown would require delegation of administrative responsibility to the group requesting the extra granularity. This architecture allows this to happen. It currently requires the group given this responsibility to administer the details of that granularity themselves.

For the latest information on Windows NT Server, check out our World Wide Web site at http://www.microsoft.com/ntserver or the Windows NT

Server Forum on the Microsoft Network (GO WORD: MSNTS).

APPENDIX A – CORPORATE
WINDOWS NT 5.0
ACTIVE DIRECTORY
NAMESPACE

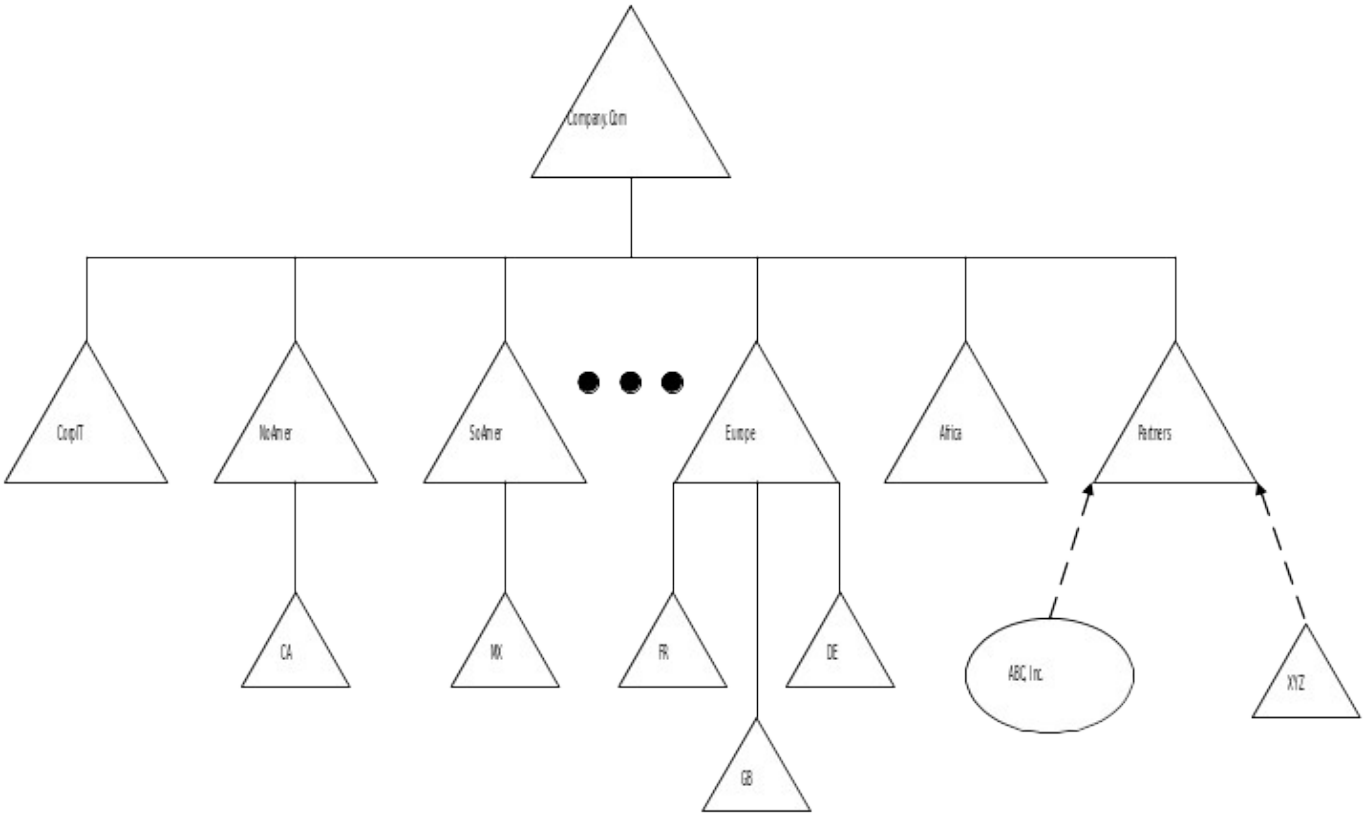

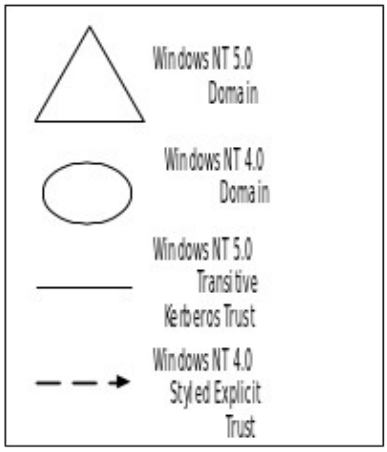

APPENDIX C – ISO 3166 TWO-CHARACTER COUNTRY CODES

APPENDIX B – SAMPLE TOP-LEVEL OU STRUCTURES

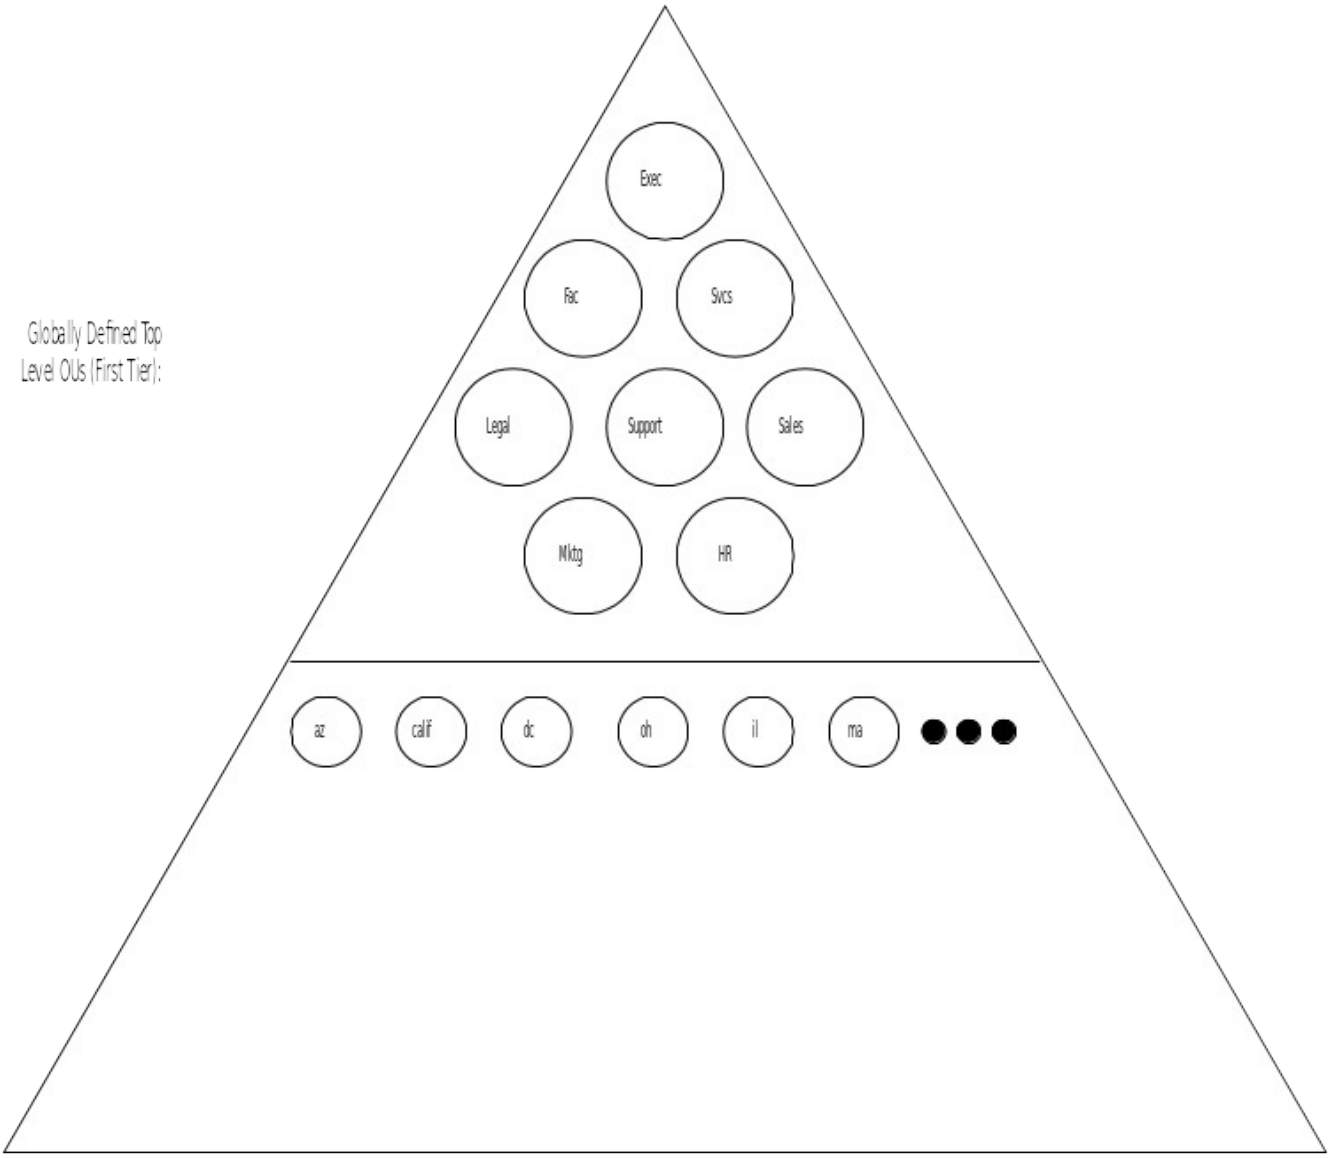

See table on following page.

ad	Andorra	cr	Costa Rica	id	Indonesia	mu	Mauritius	sk	Slovakia
ae	United Arab Emirates	cs	Czechoslovakia	ie	Ireland	mv	Maldives	sl	Sierra Leone
af	Afghanistan	cu	Cuba	il	Israel	mw	Malawi	sm	San Marino
ag	Antigua, Barbuda	cv	Cape Verde	in	India	mx	Mexico	sn	Senegal
ai	Anguilla	cx	Christmas Island	io	British Indian, Terr.	my	Malaysia	so	Somalia
al	Albania	cy	Cyprus	iq	Iraq	mz	Mozambique	sr	Suriname
am	Armenia	de	Germany, Deutschland	ir	Iran	na	Namibia	st	St. Tome and Principe
an	Netherland Antilles	dj	Djibouti	is	Iceland	nc	New Caledonia	su	Soviet Union
ao	Angola	dk	Denmark	it	Italy	ne	Niger	sv	El Salvador
aq	Antarctica	do	Dominican Republic	jm	Jamaica	nf	Norfolk Island	sy	Syria
ar	Argentina	dz	Algeria	jo	Jordan	ng	Nigeria	sz	Swaziland
as	American Samoa	ec	Ecuador	jp	Japan	ni	Nicaragua	tc	Turks &Caicos Islands
at	Austria	ee	Estonia	ke	Kenya	nl	Netherlands	td	Chad
au	Australia	eg	Egypt	kg	Kirgistan	no	Norway	tf	French Southern Terr.
aw	Aruba	eh	Western Sahara	kh	Cambodia	np	Nepal	tg	Togo
az	Azerbaidjan	es	Spain	ki	Kiribati	nr	Nauru	th	Thailand
ba	Bosnia-Herzegovina	et	Ethiopia	km	Comoros	nt	Neutral Zone	tj	Tadjikistan
bb	Barbados	fi	Finland	kp	Korea North	nu	Niue	tk	Tokelau
bd	Bangladesh	fj	Fiji	kr	Korea South	nz	New Zealand	tm	Turkmenistan
be	Belgium	fk	Falkland Island	kw	Kuwait	om	Oman	tn	Tunisia
bf	Burkina Faso	fm	Micronesia	ky	Cayman Islands	pa	Panama	to	Tonga
bg	Bulgaria	fo	Faroe Islands	kz	Kazachstan	pe	Peru	tp	East Timor
bh	Bahrain	fr	France	la	Laos	pf	Polynesia	tr	Turkey
bi	Burundi	fx	France, European Terr.	lb	Lebanon	pg	Papua New Guinea	tt	Trinidad & Tobago
bj	Benin	ga	Gabon	lc	Saint Lucia	ph	Philippines	tv	Tuvalu
bm	Bermuda	gb	Great Britain, UK	li	Liechtenstein	pk	Pakistan	tw	Taiwan
bn	Brunei Darussalam	gd	Grenada	lk	Sri Lanka	pl	Poland	tz	Tanzania
bo	Bolivia	ge	Georgia	lr	Liberia	pm	St. Pierre & Miquelon	ua	Ukraine
br	Brazil	gf	Guyana	ls	Lesotho	pn	Pitcairn	ug	Uganda
bs	Bahamas	gh	Ghana	lt	Lithuania	pr	Puerto Rico	uk	United Kingdom
bt	Buthan	gi	Gibraltar	lu	Luxembourg	pt	Portugal	um	US Minor Outlying Island
bv	Bouvet Island	gl	Greenland	lv	Latvia	pw	Palau	us	United States
bw	Botswana	gm	Gambia	ly	Libya	py	Paraguay	uy	Uruguay
by	Bielorussia	gn	Guinea	ma	Morocco	qa	Qatar	uz	Uzbekistan
bz	Belize	gp	Guadeloupe	mc	Monaco	re	Reunion	va	Vatican City State
ca	Canada	gq	Equatorial Guinea	md	Moldavia	ro	Romania	vc	St.Vincent & Grenadines
com	Commercial, Int'l. Domain Name	gr	Greece	mg	Madagascar	ru	Russian Federation	ve	Venezuela
cc	Cocos Island	gt	Guatemala	mh	Marshall Islands	rw	Rwanda	vg	Virgin Islands, British
cf	Central African Rep.	gu	Guam	ml	Mali	sa	Saudi Arabia	vi	Virgin Islands, US
cg	Congo	gw	Guinea Bissau	mm	Myanmar	sb	Solomon Islands	vn	Vietnam
ch	Switzerland	gy	Guyana	mn	Mongolia	sc	Seychelles	vu	Vanuatu
ci	Ivory Coast	hk	Hong Kong	mo	Macau	sd	Sudan	wf	Wallis &Futuna Islands
ck	Cook Islands	hm	Heard &McDonald Island	mp	Northern Mariana Island	se	Sweden	ws	Samoa
cl	Chile	hn	Honduras	mq	Martinique	sg	Singapore	ye	Yemen
cm	Cameroon	hr	Croatia	mr	Mauritania	sh	St. Helena	yu	Yugoslavia
cn	China	ht	Haiti	ms	Montserrat	si	Slovenia	za	South Africa
							Svalbard & Jan Mayen		

APPENDIX D – TWO-LETTER US POSTAL CODES

AL	Alabama	ID	Idaho	NE	Nebraska	SD	South Dakota
AK	Alaska	IL	Illinois	NV	Nevada	TN	Tennessee
AS	American Samoa	IN	Indiana	NH	New Hampshire	TX	Texas
AZ	Arizona	IA	Iowa	NJ	New Jersey	UT	Utah
AR	Arkansas	KS	Kansas	NM	New Mexico	VT	Vermont
Calif	California	KY	Kentucky	NY	New York	VA	Virginia
CZ	Canal Zone	LA	Louisiana	NC	North Carolina	VI	Virgin Islands
CO	Colorado	ME	Maine	ND	North Dakota	WA	Washington
CT	Conneticut	MD	Maryland	OH	Ohio	WV	West Virginia
DE	Deleware	MA	Massachusetts	OK	Oklahoma	WI	Wisconsin
DC	District of Columbia	MI	Michigan	OR	Oregon	WY	Wyoming
FL	Florida	MN	Minnesota	PA	Pennsylvania		
GA	Georgia	MS	Mississippi	PR	Puerto Rico		
GU	Guam	MO	Missouri	RI	Rhode Island		
HI	Hawaii	MT	Montana	SC	South Carolina		

* Note that that the actual two-letter postal code for the state of California is “CA.” We designate it as “Calif” in this paper since “CA” conflicts with the ISO 3166 code for Canada.